

What is claimed is:

1. A workload-adaptive data storage system, the data storage system receiving requests to store and retrieve data, comprising:
 - 5 a. a data storage system controller, the data storage system controller adaptively determining data organization schemes for optimizing the performance of the data storage system;
 - b. a plurality of data storage units; and
 - c. an interconnect between the data storage system controller and the plurality of data storage units.
- 10 2. The system as recited in claim 1 wherein the data storage system controller comprises:
 - a. a plurality of host adapters to connect to the storage application;
 - b. at least one processor determining the data organization scheme based on specified parameters and handling the workload on the data storage system
15 controller;
 - c. a memory buffer for data and commands;
 - d. a bus controller connecting the processor and the memory to an internal bus of the data storage system controller; and
 - e. a plurality of interconnect access units for managing communication between the
20 data storage system controller and the plurality of data storage units.
3. The system as recited in claim 2 optimizes the performance of the data storage system by determining a data storage unit for storing the data based on data storage system parameters and workload parameters.

4. The system as recited in claim 3 wherein the data storage system parameters comprise at least one of:
 - a. current storage utilization per data storage unit for each data storage unit in the data storage system;
 - 5 b. current input/output (I/O) load per data storage unit for each data storage unit in the data storage system;
 - c. I/O bandwidth per data storage unit;
 - d. number of drives currently powered on per data storage unit; and
 - e. maximum number of drives that can be powered on per data storage unit.
- 10 5. The system as recited in claim 3 wherein the workload parameters comprise at least one of:
 - a. volume size of the workload;
 - b. I/O block size;
 - c. target I/O rate for the workload; and
 - 15 d. access pattern of the workload.
6. The system as recited in claim 1 wherein the data organization schemes comprise at least one of:
 - a. RAID, the RAID scheme involving writing onto multiple disk drives simultaneously; and
 - 20 b. power managed RAID, the power managed RAID scheme involving writing onto multiple disk drives sequentially.
7. The system as recited in claim 1 wherein each data storage unit comprises:

- a. a plurality of disk systems;
 - b. a control processor handling storage requests for the disk systems;
 - c. a memory buffer for caching data and commands; and
 - d. a local interconnect connecting the plurality of disk systems to the data storage unit controller.
- 5
8. The system as recited in claim 7 wherein each data storage unit further comprises a XOR accelerator providing data protection functions;
9. The system as recited in claim 7 wherein each disk system comprises:
- a. a plurality of drives, each drive being individually controllable to power on and off,
- 10 independent of other drives; and
- b. a disk controller for managing data transfer operations.
10. A method for adaptation of data organization by a data storage system controller in a data storage system, the data storage system being supplied with workload from a storage application, the method comprising the steps of:
- a. determining workload parameters;
- 15
- b. determining data storage system parameters; and
 - c. determining a configuration of the data storage system for optimizing performance of the data storage system.
11. The method as recited in claim 10 wherein the step of determining the workload
- 20 parameters comprises the steps of:
- a. determining volume size of the workload; and
 - b. determining an I/O block size.

12. The method as recited in claim 10 wherein the step of determining the workload parameters further comprises the step of determining target I/O rate for the workload.
- 5 13. The method as recited in claim 10 wherein the step of determining the workload parameters further comprises the step of determining access pattern of the workload.
14. The method as recited in claim 10 wherein the step of determining the data storage system parameters comprises the steps of:
- 10 a. determining current storage utilization per data storage unit for each data storage unit in the data storage system;
- b. determining current I/O load per data storage unit for each data storage unit in the data storage system; and
- c. determining I/O bandwidth per data storage unit.
- 15 15. The method as recited in claim 10 wherein the step of determining the data storage system parameters further comprises the step of determining the number of drives currently powered on per data storage unit.
16. The method as recited in claim 10 wherein the step of determining the data storage system parameters further comprises the step of determining the maximum number of drives that can be powered on per data storage unit.
- 20 17. The method as recited in claim 10 wherein the step of determining the configuration of the data storage system for the workload comprises the steps of:
- a. determining target data storage unit;
- b. determining a data organization scheme for optimizing I/O performance of volumes based on the specified parameters; and

- c. determining a set of drives, in the target data storage unit, on which the volume is to be written.

18. The method as recited in claim 17 wherein the step of determining the target data storage unit is based on the lowest disk utilization.

- 5 19. The method as recited in claim 17 wherein the step of determining the data organization scheme comprises selecting RAID organization.

20. The method as recited in claim 17 wherein the step of determining the data organization scheme comprises selecting power managed RAID organization.

- 10 21. The method as recited in claim 17 wherein the step of determination of the data organization scheme is based on the power allocated to simultaneously power a maximum number of drives in the data storage unit.

22. The method as recited in claim 17 wherein the step of determination of set of drives is based on environmental considerations relating to electrical noise and heating.

- 15 23. The method as recited in claim 10 further comprising the step of allocating volumes to the best possible location within the data storage system for continuous load balancing of the storage I/O performance of all allocated storage volumes.

24. The method as recited in claim 10 further comprising the step of organizing volumes to the best I/O needs of each volume for continuous load balancing of the storage I/O performance of all allocated storage volumes.

- 20 25. A method for optimizing I/O performance and power consumption of a data storage system on receipt of a request for writing data, the I/O performance being measured in terms of the I/O rate, the data storage system comprising a plurality of data storage units and a data storage unit controller, each of the plurality of data storage units comprising multiple disk drives, the data being written on to the disk drives, the
25 method comprising the steps of:

- a. determining workload parameters, the workload parameters characterizing the workload imposed onto the data storage system
 - b. determining data storage system parameters, the data storage system parameters characterizing the data storage system;
 - 5 c. identifying target data storage units for writing the data; and
 - d. if volume of data to be written is greater than a threshold volume then
writing data using RAID configuration
else
writing data using power managed RAID configuration
- 10 26. The method as recited in claim 25 wherein identifying the target data storage units comprises the steps of:
- a. estimating possible candidate data storage units having sufficient storage capacity to allocate for writing the data; and
 - b. designating one of the possible candidate data storage units as target data
 - 15 storage unit, the target data storage unit having the best combination of storage capacity and I/O utilization.
27. The method as recited in claim 25 wherein the step of writing data using RAID configuration comprises determining the number of drives for writing the data from workload parameters.
- 20 28. The method as recited in claim 25 wherein the step of writing data using power managed RAID configuration comprises the steps of:
- a. determining the number of drives in the target data storage units for writing the data from the workload and data storage system parameters; and
 - b. sequentially writing the data to a chain of drives in the target data storage unit.